

DISCRETE LENGTH-VOLUME INEQUALITIES AND LOWER VOLUME BOUNDS IN METRIC SPACES

KYLE KINNEBERG

ABSTRACT. A theorem of W. Derrick ensures that the volume of any Riemannian cube $([0, 1]^n, g)$ is bounded below by the product of the distances between opposite codimension-1 faces. In this paper, we establish a discrete analog of Derrick's inequality for weighted open covers of the cube $[0, 1]^n$, which is motivated by a question about lower volume bounds in metric spaces. Our main theorem generalizes a previous result of the author in [13], which gave a combinatorial version of Derrick's inequality and was used in the analysis of boundaries of hyperbolic groups. As an application, we answer a question of Y. Burago and V. Zalgaller about length-volume inequalities for pseudometrics on the unit cube.

1. INTRODUCTION

There is a deep and well-studied relationship in metric geometry between the volume of a space and the lengths of curves that, in some way, generate it. An early example of such a relationship is due to K. Loewner (unpublished, but see [18] for a discussion) and deals with conformal structures on the torus \mathbb{T}^2 .

Theorem 1.1. *Let (\mathbb{T}^2, g) be the 2-dimensional torus, equipped with a Riemannian metric g , let $\ell(g)$ denote the infimal length of a closed curve on \mathbb{T}^2 that is not homotopically trivial, and let $\text{Vol}(g)$ denote the volume of \mathbb{T}^2 with respect to the metric g . Then $\text{Vol}(g) \geq \frac{\sqrt{3}}{2}\ell(g)^2$, and equality holds if and only if (\mathbb{T}^2, g) is isometric to the flat torus \mathbb{R}^2/Λ , where Λ is the lattice generated by $(1, 0)$ and $(1/2, \sqrt{3}/2)$.*

Loewner's inequality is only the beginning of a very rich body of work that has sought to understand similar phenomena for more general spaces and in dimensions greater than two. We refer to [8, Chapter 4] for a broad survey of methods and results in this area. Of particular interest to us is the following theorem, originally proved by W. Derrick [5, Theorem 3.4]. Here we state it in the form cited in [8].

Theorem 1.2. *Let $([0, 1]^n, g)$ be the n -dimensional unit cube, equipped with a Riemannian metric g . Let F_k and F'_k denote the pairs of opposite codimension-1 faces of $[0, 1]^n$, for $1 \leq k \leq n$, and let d_k be the distance between F_k and F'_k with respect to the metric determined by g . Then $\text{Vol}(g) \geq d_1 \cdots d_n$.*

Let us outline the main ideas in the proof of Derrick's inequality; we will see them reappear in combinatorial form later in this paper. For each k , let $f_k: [0, 1]^n \rightarrow \mathbb{R}$ be defined by $f_k(x) = \text{dist}_g(x, F_k)$, where dist_g denotes the distance with respect to the metric determined by g . It is clear that $f_k(F_k) = \{0\}$ for each k , and the

Date: February 24, 2016.

2010 Mathematics Subject Classification. Primary: 52C17; Secondary: 30L99.

definition of d_k ensures that $f_k(F'_k) \subset [d_k, \infty)$. Moreover, each f_k is 1-Lipschitz. For simplicity, let us assume that f_k is smooth (in reality, one would approximate f_k by smooth functions that are $(1 + \epsilon)$ -Lipschitz), so the differential satisfies $\|df_k\| \leq 1$. Now define $f: [0, 1]^n \rightarrow \mathbb{R}^n$ by $f(x) = (f_1(x), \dots, f_n(x))$. If $\omega = dy_1 \wedge \dots \wedge dy_n$ denotes the canonical volume form on \mathbb{R}^n , then its pull-back is $f^*\omega = df_1 \wedge \dots \wedge df_n$. Hadamard's inequality then gives

$$\|f^*\omega\| \leq \|df_1\| \cdots \|df_n\| \leq 1.$$

Letting ν denote the Riemannian volume form on $([0, 1]^n, g)$, we have

$$(1.1) \quad \text{Vol}(g) \geq \int_{[0, 1]^n} \|f^*\omega\| d\nu \geq \left| \int_{[0, 1]^n} f^*\omega \right| \geq \int_{f([0, 1]^n)} \omega.$$

As $f_k(F_k) = \{0\}$ and $f_k(F'_k) \subset [d_k, \infty)$ for each k , standard topological arguments ensure that the rectangle $[0, d_1] \times \dots \times [0, d_n]$ is contained in $f([0, 1]^n)$ (see Lemma 2.2 below). Consequently, the product $d_1 \cdots d_n$ is a lower bound for the right-hand side of (1.1), and this gives the desired inequality.

In [13, Section 4.3], the present author proved a combinatorial version of Derrick's inequality for open covers of the unit cube $[0, 1]^n$. This result was used in the construction of a metric with certain regularity properties on the boundary of a Gromov hyperbolic metric space. The set-up was as follows. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $[0, 1]^n$, and again let F_k and F'_k denote the pairs of opposite codimension-1 faces. We say that U_{i_1}, \dots, U_{i_m} is a *chain* if $U_{i_j} \cap U_{i_{j+1}} \neq \emptyset$ for each j . Moreover, such a chain is said to connect two sets A and B if $U_{i_1} \cap A \neq \emptyset$ and $U_{i_m} \cap B \neq \emptyset$.

Theorem 1.3 ([13, Proposition 4.4]). *Let d_k denote the smallest number of sets U_i in a chain that connects F_k and F'_k . Then $\#\mathcal{U} \geq d_1 \cdots d_n$.*

Note that although this result is analogous to Derrick's theorem, it does not parallel the Riemannian inequality. Indeed, the sets U_i in \mathcal{U} are essentially treated as if they all had diameter 1. The primary purpose of this paper is to extend the preceding result to a weighted version, which is closer to Theorem 1.2 and also generalizes Theorem 1.3.

To this end, let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $[0, 1]^n$ as before, and let $w: I \rightarrow [0, \infty)$ be a corresponding weight function. One should think of $w(i)$ as the weight associated to the set U_i . Together, \mathcal{U} and w give us a discrete notion of distance on $[0, 1]^n$. Namely, we define

$$\text{dist}_w(A, B) = \inf \left\{ \sum_{j=1}^m w(i_j) : \begin{array}{l} U_{i_1}, \dots, U_{i_m} \text{ is a chain} \\ \text{that connects } A \text{ and } B \end{array} \right\}.$$

The sum $w(i_1) + \dots + w(i_m)$ is said to be the length of the corresponding chain. By path connectedness of $[0, 1]^n$ and compactness of paths, it is easy to see that any two points in the cube can be connected by a chain. In particular, $\text{dist}_w(A, B)$ is finite for $A, B \subset [0, 1]^n$. We should note that a chain might be disconnected topologically, as we have made no assumption on the connectedness of the sets in \mathcal{U} . Our main theorem in this paper is the following.

Theorem 1.4. *Let \mathcal{U} be an open cover of $[0, 1]^n$, let w be a corresponding weight function, and let $d_k = \text{dist}_w(F_k, F'_k)$ for each $1 \leq k \leq n$. Then*

$$\sum_{i \in I} w(i)^n \geq d_1 \cdots d_n.$$

In fact, we will prove a more general version of this inequality that has more in common with the results of Derrick in [4] and [6]: we allow the discrete distance between F_k and F'_k to be taken with respect to (possibly) different weight functions for different values of k .

Theorem 1.5. *Let \mathcal{U} be an open cover of $[0, 1]^n$, and let w_k be associated weight functions for $1 \leq k \leq n$. If $d_k = \text{dist}_{w_k}(F_k, F'_k)$ for each k , then*

$$\sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right) \geq d_1 \cdots d_n.$$

It is clear that Theorem 1.4 follows immediately from Theorem 1.5 by setting $w_k = w$ for each k .

As a corollary to Theorem 1.4, we will easily obtain lower *Hausdorff content* bounds for continuous images of $[0, 1]^n$ in arbitrary metric spaces. Recall that if (X, d) is a metric space, the Q -dimensional Hausdorff content of a compact subset $E \subset X$ is defined to be

$$\mathcal{H}_Q^\infty(E) = \inf \left\{ \sum_{i \in I} (\text{diam } U_i)^Q : \{U_i\}_{i \in I} \text{ is an open cover of } E \right\}.$$

We will show the following bound.

Corollary 1.6. *Let $g: [0, 1]^n \rightarrow X$ be continuous. Then*

$$\mathcal{H}_n^\infty(g([0, 1]^n)) \geq \prod_{k=1}^n \text{dist}(g(F_k), g(F'_k)),$$

where dist denotes the metric distance between sets in X .

For example, if $([0, 1]^n, d)$ is the unit cube equipped with an arbitrary metric whose topology coincides with the Euclidean topology, then one can apply this inequality to the identity function. We remark here that our definition of Hausdorff content does not include the normalizing multiplicative factor $\text{Vol}(B^n)2^{-n}$, where B^n is the unit ball in \mathbb{R}^n and $\text{Vol}(\cdot)$ is n -dimensional Lebesgue measure, as is standard in geometric measure theory and which appears in the results of Derrick. In this sense, Corollary 1.6 does not recover Derrick's inequality. On the other hand, Corollary 1.6 is sharp in the metric category, and it can also be used to answer a question posed by Y. Burago and V. Zalgaller in [3, p. 296] concerning pseudometrics on $[0, 1]^n$. This is discussed further in Section 4.

More generally, Corollary 1.6 provides a method of verifying lower volume bounds in fairly general classes of metric spaces. Namely, if one can find continuous images of cubes in a metric space X whose sides are well-separated, then X must satisfy some corresponding lower volume bounds. This leads us to study metric spaces that “admit fat cubes.” It turns out that some standard types of connectivity properties (that are common, for example, in the study of boundaries of hyperbolic groups) ensure that the metric space admits fat squares.

The outline of the paper is as follows. In Section 2, we will introduce some techniques that appear frequently: partitions of unity, the nerve of open covers, and the topological non-degeneracy lemma that we used in the above discussion on Derrick's inequality. Section 3 will be devoted entirely to the proof of Theorem 1.5. In Section 4, we take up the topic of lower Hausdorff content bounds in metric spaces, and the final section extends these considerations to the study of metric surfaces.

Acknowledgments. This paper is part of the author's Ph.D. thesis at UCLA. He thanks his advisor, Mario Bonk, for many helpful discussions and recommendations about this project. He also thanks Andrey Mishchenko for bringing the (crucial) paper [19] to his attention and John Mackay for helpful feedback.

The author gratefully acknowledges partial support from NSF grant DMS-1162471 during his work on this project.

2. NOTATION AND PRELIMINARIES

Let (X, d) be a metric space. As is standard, we will use $B(x, r)$ to denote the open ball centered at $x \in X$ with radius $r > 0$. For subsets $A, B \subset X$, we let

$$\text{dist}(A, B) = \inf\{d(x, y) : x \in A \text{ and } y \in B\}$$

be the distance between A and B . In the case that $A = \{x\}$ we abuse notation and simply write $\text{dist}(x, B)$. We also let

$$\text{diam } A = \sup\{d(x, y) : x, y \in A\}$$

be the diameter of the subset A . Following common notation, we use $\text{int}(A)$ and \overline{A} to denote, respectively, the interior and closure of A (the ambient space for the closure operation will be understood from context). Also, we let $\partial A = \overline{A} \setminus \text{int}(A)$ denote the boundary of A . Lastly, for $\epsilon > 0$, let

$$N_\epsilon(A) = \{x \in X : \text{dist}(x, A) < \epsilon\}$$

be the open ϵ -neighborhood of A .

Suppose that (X, d) is compact, and let $\mathcal{U} = \{U_i\}_{i \in I}$ be a finite open cover. By the Lebesgue lemma, there is a positive constant $\delta > 0$ such that each ball $B(x, 2\delta)$ lies entirely in some set U_i . Let

$$f_i(x) = \min\left\{1, \frac{1}{\delta} \text{dist}(x, N_\delta(X \setminus U_i))\right\},$$

which is a $1/\delta$ -Lipschitz function with values in $[0, 1]$ and whose support is contained in U_i . Moreover, for each $x \in X$, we have

$$f(x) := \sum_{i \in I} f_i(x) \geq 1,$$

as $B(x, 2\delta) \subset U_i$ for some i . If $N = \max\{\sum_i \chi_{U_i}(x) : x \in X\}$, then also $f(x) \leq N$ for all x . Define $\phi_i(x) = f_i(x)/f(x)$ so that the following properties hold:

- (i) ϕ_i is $(2N + 1)/\delta$ -Lipschitz with support contained in U_i ;
- (ii) $0 \leq \phi_i(x) \leq 1$ for all $x \in X$;
- (iii) $\sum_i \phi_i(x) = 1$ for all $x \in X$.

The family $\{\phi_i\}$ is therefore a $(2N + 1)/\delta$ -Lipschitz *partition of unity* subordinate to \mathcal{U} [10, Chapter 2]. Partitions of unity are useful in general metric settings to produce a type of proxy for linear structure.

There is a canonical way to associate a simplicial complex to the cover \mathcal{U} whose combinatorics mimics the combinatorics of \mathcal{U} . Often, this simplicial complex is defined as an abstract complex that encodes the intersections among the sets in \mathcal{U} . We prefer to work with a geometric realization of this complex in Euclidean space. For ease, then, let us index the collection $\{U_i\}$ by the integers $1, \dots, M$, and let e_i be the i -th standard basis vector in \mathbb{R}^M .

Definition 2.1. The nerve of \mathcal{U} , denoted by $\text{Ner}(\mathcal{U})$, is

$$\text{Ner}(\mathcal{U}) = \bigcup \{ \text{conv}(e_{i_0}, \dots, e_{i_m}) : U_{i_0} \cap \dots \cap U_{i_m} \neq \emptyset \},$$

where the union runs over collections of sets in \mathcal{U} that have non-empty intersection.

Here, and in general, we use $\text{conv}(A)$ to denote the convex hull of a set $A \subset \mathbb{R}^M$. When $A = \{a_0, \dots, a_m\}$ is a finite set, we can express

$$(2.1) \quad \text{conv}(A) = \left\{ \sum_{i=0}^m \lambda_i a_i : \lambda_i \geq 0 \text{ and } \lambda_0 + \dots + \lambda_m = 1 \right\},$$

and if $m \leq M$, then this is a (possibly degenerate) m -dimensional simplex in \mathbb{R}^M . Thus, the simplex spanned by e_{i_0}, \dots, e_{i_m} in \mathbb{R}^M is in the nerve of \mathcal{U} if, and only if, the corresponding sets U_{i_0}, \dots, U_{i_m} have a common intersection.

The partition of unity $\{\phi_i\}_{i \in I}$ allows us to map X naturally to $\text{Ner}(\mathcal{U})$. Namely, define $\phi: X \rightarrow \text{Ner}(\mathcal{U})$ by

$$(2.2) \quad \phi(x) = \sum_{i \in I} \phi_i(x) e_i, \quad x \in X,$$

and note that ϕ is continuous. In fact, as each ϕ_i is Lipschitz, the map ϕ will be Lipschitz as well. The fact that $\phi(x) \in \text{Ner}(\mathcal{U})$ follows immediately from the definition of the nerve, the characterization in (2.1), and the properties (ii) and (iii) above.

It will be useful for us later to subdivide the simplices in the nerve without changing $\text{Ner}(\mathcal{U})$ as a set in \mathbb{R}^M . The *barycentric subdivision* allows us to do this in a canonical way. Once again, we work with a geometric realization of the relevant complexes.

Let $S \subset \mathbb{R}^M$ be a simplicial complex whose simplices are convex hulls of the standard basis vectors e_i . For each collection $\{e_{i_0}, \dots, e_{i_m}\}$ of vertices which generate a simplex in S , we define its *barycenter* to be the point

$$\text{bc}(e_{i_0}, \dots, e_{i_m}) = \frac{1}{m+1} (e_{i_0} + \dots + e_{i_m}).$$

The subdivision proceeds inductively, by dimension, on the simplices in S . Intuitively, we may think about it in the following way. First, subdivide each edge by adding a vertex at $\text{bc}(e_i, e_j)$ whenever $\text{conv}(e_i, e_j)$ is in S . Second, subdivide each 2-dimensional simplex by adding a vertex at $\text{bc}(e_i, e_j, e_k)$ whenever $\text{conv}(e_i, e_j, e_k)$ is in S , and then add edges from $\text{bc}(e_i, e_j, e_k)$ to each vertex on the boundary of $\text{conv}(e_i, e_j, e_k)$ (these vertices may come from S itself or from the first step in the subdivision). Continue in the same way, until each simplex in S has been subdivided. For further reference, see [9, pp. 119–120].

The resulting simplicial complex is called the first barycentric subdivision of S . Observe that the geometric realizations of the complexes we obtain throughout this process, including in the final step, coincide with S as sets in \mathbb{R}^M . We will, however, use S_b to denote the geometric realization of this new complex to emphasize the fact

that we have a refined simplicial structure. The following fact will be important in later arguments. If $\text{conv}(e_{i_0}, \dots, e_{i_m})$ is an m -dimensional simplex in S , then after the barycentric subdivision, it is a union of m -dimensional simplices in S_b with geometric form

$$(2.3) \quad \text{conv}(p_0, \dots, p_m),$$

where $p_j = \text{bc}(e_{\sigma(i_0)}, \dots, e_{\sigma(i_j)})$ for each $0 \leq j \leq m$, and σ is a permutation of the indices i_0, \dots, i_m . In particular, $p_0 = e_{\sigma(i_0)}$.

Before concluding this section, let us record a topological lemma that will be useful later on. Let P be a compact, convex set in \mathbb{R}^n . We say that a closed half-space $H \subset \mathbb{R}^n$ *supports* P if $P \cap H$ is non-empty and is contained in ∂H .

Lemma 2.2. *Let $P \subset \mathbb{R}^n$ be compact and convex, with non-empty interior. Suppose that $f: P \rightarrow \mathbb{R}^n$ is continuous, and for each $x \in \partial P$ there is a closed half-space H which supports P , with $x, f(x) \in H$. Then $P \subset f(P)$.*

Proof. As $f(P)$ is compact and P is the closure of $\text{int}(P)$, it suffices to show that $\text{int}(P) \subset f(P)$. Aiming for a contradiction, let us suppose that there is $y \in \text{int}(P) \setminus f(P)$. By translation, we may assume that $y = 0$.

The compactness and convexity of P imply that for each $x \in \mathbb{R}^n \setminus \{0\}$, the ray from 0 through x intersects ∂P in exactly one point, which we denote by $p(x)$. We claim that the map $p: \mathbb{R}^n \setminus \{0\} \rightarrow \partial P$ is continuous. To verify this, fix $\epsilon > 0$ small enough that $\overline{B}(0, \epsilon) \subset \text{int}(P)$, and let $\pi: \mathbb{R}^n \setminus \{0\} \rightarrow \partial B(0, \epsilon)$ be the canonical projection onto the sphere of radius ϵ . It is clear that π is continuous. Observe also that $p(x) = p(\pi(x))$ for all $x \in \mathbb{R}^n \setminus \{0\}$, so it suffices to show that p , restricted to $\partial B(0, \epsilon)$, is continuous. For this, we note that $p|_{\partial B(0, \epsilon)}$ is the inverse map of $\pi|_{\partial P}$, the latter of which is a continuous bijection from the compact set ∂P to $\partial B(0, \epsilon)$. Consequently, it is a homeomorphism, so its inverse $p|_{\partial B(0, \epsilon)}$ is also continuous.

Consider the map $g: P \rightarrow \partial P$ defined by $g(x) = p(-f(x))$. This is continuous by the assumption that $0 \notin f(P)$. We claim that g has no fixed point. Indeed, if $g(x) = x$, then necessarily $x \in \partial P$. By hypothesis, there is a half-space H of \mathbb{R}^n which supports P and has $x, f(x) \in H$. As 0 lies in the complement of H , we know that $-f(x)$ is in the complement of H as well. The point $p(-f(x))$ lies on the segment joining 0 and $-f(x)$, and so it also fails to lie in H . This, however, contradicts the fact that $g(x) = x \in H$, so g can have no fixed points. The existence of such a map g contradicts the Brouwer fixed point theorem: any continuous map from a compact, convex set in \mathbb{R}^n to itself has a fixed point. Thus, we obtain $\text{int}(P) \subset f(P)$, as desired. \square

3. A TOPOLOGICAL LENGTH-VOLUME INEQUALITY FOR CUBES

Our primary goal in this section is to prove Theorem 1.5; recall that Theorem 1.4 follows immediately from it. Before starting the proof, we should remark that our methods are heavily informed by the ideas in O. Schramm's paper on square tilings of rectangles [19]. In fact, there are many similarities between our proof of Theorem 1.5 and the proof of [19, Theorem 5.1].

Let $[0, 1]^n$ be the standard Euclidean unit cube of dimension $n \geq 1$. We will use F_k and F'_k , for $1 \leq k \leq n$, to denote the pairs of opposite codimension-1 faces of the unit cube:

$$F_k = [0, 1]^n \cap \pi_k^{-1}(\{0\}) \quad \text{and} \quad F'_k = [0, 1]^n \cap \pi_k^{-1}(\{1\}),$$

where $\pi_k: \mathbb{R}^n \rightarrow \mathbb{R}$ is the projection to the k -th coordinate axis.

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $[0, 1]^n$, and let $w_k: I \rightarrow [0, \infty)$ be corresponding weight functions for $1 \leq k \leq n$. Here, and in what follows, an open cover of $[0, 1]^n$ will always mean that the sets are open in the relative topology on $[0, 1]^n$, unless otherwise explicitly stated. Using the notation from Section 1, we let $d_k = \text{dist}_{w_k}(F_k, F'_k)$. To prove Theorem 1.5, we must show that

$$(3.1) \quad \sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right) \geq d_1 \cdots d_n.$$

To obtain this inequality, we first work under an additional technical assumption on the open cover \mathcal{U} . Namely, if there exists $i \in I$ for which $U_i \cap F_k \neq \emptyset$ and $U_i \cap F'_k \neq \emptyset$ for some k , then we say that \mathcal{U} is *spanning*. We will verify the desired inequality first in the case that \mathcal{U} is *non-spanning*. After doing this, we will treat the general case by modifying slightly the open cover under consideration. Let us state the intermediate result as a separate proposition.

Proposition 3.1. *The inequality in (3.1) holds under the hypothesis that the cover \mathcal{U} is non-spanning.*

Proof. We may, of course, assume that $d_k > 0$ for each k ; otherwise, the desired inequality is trivial. We may also assume that the cover \mathcal{U} is finite. Indeed, compactness guarantees that any cover \mathcal{U} of the cube contains a finite sub-cover. Removing the “redundant” sets from this collection does not increase the left-hand side of the desired inequality and also does not decrease the distances d_k . Our proof will now proceed in several steps, which we explicitly indicate.

Step 1: Associate a rectangle to each U_i . For $i \in I$, define

$$d_k(i) = \begin{cases} 0 & \text{if } U_i \cap F_k \neq \emptyset, \\ \text{dist}_{w_k}(F_k, U_i) & \text{otherwise} \end{cases}$$

for $1 \leq k \leq n$. Of course, we have $d_k(i) \geq 0$ for each i and k . Also note that, by finiteness of the cover \mathcal{U} , the infimum in the definition of discrete distances can be replaced by a minimum. More importantly, however, is the fact that if $U_i \cap F'_k \neq \emptyset$, then $d_k \leq d_k(i) + w_k(i)$. This follows immediately from the relevant definitions.

Now define

$$R_i = \prod_{k=1}^n [d_k(i), d_k(i) + w_k(i)],$$

which is an n -dimensional rectangle with side lengths $w_k(i)$. To simplify notation, we let

$$I_k(i) = \pi_k(R_i) = [d_k(i), d_k(i) + w_k(i)],$$

so that $R_i = \prod_k I_k(i)$.

We will use R_i as a sort of proxy for the set U_i . It will therefore be important that the combinatorics of the rectangles $\{R_i\}_{i \in I}$ mimic those of the sets $\{U_i\}_{i \in I}$, in the following sense.

Claim 1. If $U_i \cap U_j \neq \emptyset$, then $R_i \cap R_j \neq \emptyset$.

Proof of claim. We simply need to show that $I_k(i) \cap I_k(j) \neq \emptyset$ for each k . To this end, fix k and without loss of generality, assume that $d_k(i) \leq d_k(j)$. We claim that $d_k(j) \leq d_k(i) + w_k(i)$. Indeed, there is a chain U_{i_1}, \dots, U_{i_m} that connects F_k and U_i of length $d_k(i) = \text{dist}_{w_k}(F_k, U_i)$; in case $U_i \cap F_k \neq \emptyset$, this is simply the empty

chain. As $U_i \cap U_j \neq \emptyset$, the augmented chain $U_{i_1}, \dots, U_{i_m}, U_i$ connects F_k and U_j . Thus, $d_k(j) \leq d_k(i) + w_k(i)$, which, along with the assumption that $d_k(i) \leq d_k(j)$, immediately gives $I_k(i) \cap I_k(j) \neq \emptyset$. \square

We should remark here that the converse need not hold; there are many configurations, in fact, for which two rectangles intersect even though the corresponding open sets are disjoint. Claim 1 easily gives the following.

Claim 2. If $U_{i_0} \cap \dots \cap U_{i_m} \neq \emptyset$, then $R_{i_0} \cap \dots \cap R_{i_m} \neq \emptyset$.

Proof of claim. As before, it suffices to show that $I_k(i_0) \cap \dots \cap I_k(i_m) \neq \emptyset$ for each k . From the previous claim, we know that $R_{i_j} \cap R_{i_{j'}} \neq \emptyset$ for any pair j, j' ; in particular, $I_k(i_j) \cap I_k(i_{j'}) \neq \emptyset$. Thus, the $m+1$ intervals $I_k(i_0), \dots, I_k(i_m)$ have pairwise non-empty intersections. This implies that $I_k(i_0), \dots, I_k(i_m)$ have a point of common intersection: indeed, the maximum among their left endpoints is at most the minimum among their right endpoints. \square

Now that we have established a correspondence between the combinatorics of $\{U_i\}_{i \in I}$ and $\{R_i\}_{i \in I}$, we wish to map the unit cube $[0, 1]^n$ continuously into $\bigcup_i R_i$. To ensure that the image lies within the union of the rectangles, it is technically convenient to pass through the nerve of the cover \mathcal{U} . Recall from Section 2 that by enumerating $\mathcal{U} = \{U_1, \dots, U_M\}$, we can express

$$\text{Ner}(\mathcal{U}) = \bigcup \{ \text{conv}(e_{i_0}, \dots, e_{i_m}) : U_{i_0} \cap \dots \cap U_{i_m} \neq \emptyset \}.$$

The associated partition of unity $\{\phi_i\}$ subordinate to \mathcal{U} , which we constructed in the previous section, gives the continuous map

$$\phi: [0, 1]^n \rightarrow \text{Ner}(\mathcal{U})$$

that was introduced in (2.2).

Step 2: Map $\text{Ner}(\mathcal{U})$ into $\bigcup_i R_i$. In order to map $\text{Ner}(\mathcal{U})$ into the union of the rectangles R_i , we will pass to the first barycentric subdivision of the nerve and then define our map simplicially. For ease, we use S to denote the complex $\text{Ner}(\mathcal{U})$ and, consistent with earlier notation, the complex obtained after the subdivision will be denoted by S_b . As sets in \mathbb{R}^M , the complexes S and S_b coincide; moreover, each vertex in S_b arises as the barycenter of a simplex in S .

To define $\psi: S_b \rightarrow \bigcup_i R_i$, let us first determine where it sends the vertices. Fix such a vertex p , so that $p = \text{bc}(e_{i_0}, \dots, e_{i_m})$ for some simplex, $\text{conv}(e_{i_0}, \dots, e_{i_m})$, in the nerve S . Note that the choice of e_{i_0}, \dots, e_{i_m} is uniquely determined by p , up to order. Then, as $U_{i_0} \cap \dots \cap U_{i_m} \neq \emptyset$, Claim 2 above guarantees that

$$R_{i_0} \cap \dots \cap R_{i_m} \neq \emptyset.$$

We want to send p to a point z_p in this intersection, but we must be careful how to choose it. Recall that

$$\begin{aligned} R_{i_0} \cap \dots \cap R_{i_m} &= \left(\prod_{k=1}^n I_k(i_0) \right) \cap \dots \cap \left(\prod_{k=1}^n I_k(i_m) \right) \\ (3.2) \quad &= \prod_{k=1}^n I_k(i_0) \cap \dots \cap I_k(i_m), \end{aligned}$$

so choosing z_p in $R_{i_0} \cap \dots \cap R_{i_m}$ amounts to choosing each coordinate $\pi_k(z_p)$ in the interval

$$[a_k, b_k] := I_k(i_0) \cap \dots \cap I_k(i_m).$$

We do this according to the following rule. If $U_{i_j} \cap F_k \neq \emptyset$ for each j , then we choose $\pi_k(z_p) = a_k$; observe that in this case, $a_k = 0$. Otherwise, we choose $\pi_k(z_p) = b_k$.

Let $\psi(p) = z_p$ be as above for the vertices p of S_b . Extend ψ to be affine on each simplex in S_b so that $\psi: S_b \rightarrow \mathbb{R}^n$ is continuous. We claim that the image is contained in $\bigcup_i R_i$. To see this, first observe that we may express S_b as a union of simplices Δ that are obtained by subdividing a simplex in S of the same dimension. By (2.3) in the previous section, such simplices have geometric form

$$\Delta = \text{conv}(p_0, \dots, p_m),$$

where $p_j = \text{bc}(e_{i_0}, \dots, e_{i_j})$ and $\text{conv}(e_{i_0}, \dots, e_{i_m})$ is a simplex in S . Consequently,

$$\psi(\Delta) = \text{conv}(\psi(p_0), \dots, \psi(p_m)) = \text{conv}(z_{p_0}, \dots, z_{p_m}).$$

The choice of z_{p_j} guarantees that

$$z_{p_j} \in R_{i_0} \cap \dots \cap R_{i_j} \subset R_{i_0}$$

for each $j = 0, \dots, m$. As R_{i_0} is convex, we find $\psi(\Delta) \subset R_{i_0}$.

Step 3: Map $[0, 1]^n$ into $\bigcup_i R_i$. We now want to compose $\phi: [0, 1]^n \rightarrow S$ and $\psi: S_b \rightarrow \bigcup_i R_i$ to obtain a map from the unit cube into the collection of rectangles. Recall that the complexes S and S_b coincide as sets in \mathbb{R}^M , so we can define

$$f = \psi \circ \phi: [0, 1]^n \rightarrow \bigcup_{i \in I} R_i,$$

which is continuous. Our goal now is to show that the image of f contains the n -dimensional rectangle $R = \prod_k [0, d_k]$.

The main claim that we must establish toward this end is that

$$f(F_k) \subset \pi_k^{-1}(\{0\}) \quad \text{and} \quad f(F'_k) \subset \pi_k^{-1}([d_k, \infty)).$$

From here, Lemma 2.2 almost immediately implies that $R \subset f([0, 1]^n)$. To begin, let $x \in F_k \cup F'_k$, and let U_{i_0}, \dots, U_{i_m} be the sets in \mathcal{U} that contain x . Then

$$\phi(x) = \sum_{j=0}^m \phi_{i_j}(x) e_{i_j},$$

and $x \in \bigcap_j U_{i_j}$ implies that $\text{conv}(e_{i_0}, \dots, e_{i_m})$ is a simplex in $\text{Ner}(\mathcal{U})$. Also observe that if $x \in F_k$, then $U_{i_j} \cap F_k \neq \emptyset$ for each j ; similarly, if $x \in F'_k$, then $U_{i_j} \cap F'_k \neq \emptyset$ for each j .

As $\phi(x) \in \text{conv}(e_{i_0}, \dots, e_{i_m})$, after the barycentric subdivision, we know that

$$\phi(x) \in \Delta = \text{conv}(p_0, \dots, p_m),$$

where $p_j = \text{bc}(e_{i_0}, \dots, e_{i_j})$ for each j (without loss of generality, we may re-order the indices so that σ is the identity permutation). Consequently, $f(x) = \psi(\phi(x))$ is a convex combination of the points $\psi(p_j) = z_{p_j}$. It therefore suffices to show that $\pi_k(z_{p_j}) = 0$ for each j if $x \in F_k$, and that $\pi_k(z_{p_j}) \geq d_k$ for each j if $x \in F'_k$.

In the former case, we have $U_{i_j} \cap F_k \neq \emptyset$, so that $d_k(i_j) = 0$ for each j . Consequently, we know that $I_k(i_j) = [0, w_k(i_j)]$, so

$$I_k(i_0) \cap \dots \cap I_k(i_j) = [0, b_k(j)]$$

for some $b_k(j) \geq 0$. The choice of z_{p_j} then guarantees that $\pi_k(z_{p_j}) = 0$ for each j .

In the latter case, we have $U_{i_j} \cap F'_k \neq \emptyset$, so that

$$d_k \leq d_k(i_j) + w_k(i_j)$$

for each j . In particular,

$$I_k(i_0) \cap \cdots \cap I_k(i_j) = [a_k(j), b_k(j)]$$

for some $a_k(j) \geq 0$ and $b_k(j) \geq d_k$. As \mathcal{U} is non-spanning, we know that $U_{i_j} \cap F_k = \emptyset$ for each j . By the choice of z_{p_j} , we therefore have

$$\pi_k(z_{p_j}) = b_k(j) \geq d_k$$

for each j , as desired.

It is now straightforward to conclude the proof using Lemma 2.2. Namely, let H_k be the half-space $\pi_k^{-1}([0, \infty))$, and let H'_k be the half-space $\pi_k^{-1}((-\infty, d_k])$, so that

$$R = \prod_{k=1}^n [0, d_k] = \left(\bigcap_{k=1}^n H_k \right) \cap \left(\bigcap_{k=1}^n H'_k \right).$$

Let $G_k = R \cap \partial H_k$ and $G'_k = R \cap \partial H'_k$ be the faces of R corresponding to F_k and F'_k , respectively, and let $g: R \rightarrow [0, 1]^n$ be the linear map with $g(G_k) = F_k$ and $g(G'_k) = F'_k$. We showed above that $f(F_k) \subset \mathbb{R}^n \setminus \text{int}(H_k)$ and $f(F'_k) \subset \mathbb{R}^n \setminus \text{int}(H'_k)$, so the composition $f \circ g: R \rightarrow \mathbb{R}^n$ has

$$f \circ g(G_k) \subset \mathbb{R}^n \setminus \text{int}(H_k) \quad \text{and} \quad f \circ g(G'_k) \subset \mathbb{R}^n \setminus \text{int}(H'_k).$$

Lemma 2.2 then guarantees that $R \subset f \circ g(R) = f([0, 1]^n)$. As $f([0, 1]^n) \subset \bigcup_i R_i$, volume considerations immediately give

$$d_1 \cdots d_n = \text{Vol}(R) \leq \sum_{i \in I} \text{Vol}(R_i) = \sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right),$$

as desired. \square

It is not difficult now to prove Theorem 1.5; we only need to argue that the non-spanning assumption in Proposition 3.1 is not necessary.

Proof of Theorem 1.5. Let \mathcal{U} be an open cover of $[0, 1]^n$, let w_k be associated weight functions, and let $d_k = \text{dist}_{w_k}(F_k, F'_k)$ be the corresponding distances, as in the statement of the theorem. Just as in the beginning of the proof of Proposition 3.1, it suffices to assume that \mathcal{U} is finite. Our goal is to modify the cover and the weights slightly in order to obtain a new cover to which we can apply Proposition 3.1. We will do this in such a way that the “volume” and the “lengths” associated to the new cover are very close to the original quantities. We will perform this modification in multiple steps.

First, we wish to modify \mathcal{U} to obtain an open cover of the cube so that any two sets either intersect or have strictly positive distance from each other. To this end, let $\delta_1 > 0$ be small enough that for each $x \in [0, 1]^n$, the ball $B(x, \delta_1)$ is entirely contained in some U_i . Also, let $\delta_2 > 0$ be small enough so that whenever $U_i \cap U_j \neq \emptyset$, there is some point z with $B(z, \delta_2) \subset U_i \cap U_j$. Similarly, let $\delta_3 > 0$ be small enough so that whenever $U_i \cap F_k \neq \emptyset$, there is $z \in F_k$ with $B(z, \delta_3) \subset U_i$, and whenever $U_i \cap F'_k \neq \emptyset$, there is $z \in F'_k$ with $B(z, \delta_3) \subset U_i$. Now define $\delta = \min\{\delta_1, \delta_2, \delta_3\}$.

For each i , let $A_i = \{j : U_i \cap U_j = \emptyset\}$, $B_i = \{k : U_i \cap F_k = \emptyset\}$, and $B'_i = \{k : U_i \cap F'_k = \emptyset\}$. We then form the sets

$$\tilde{U}_i = U_i \setminus \left(\bigcup_{j \in A_i} \overline{N}_{\delta/2}(U_j) \cup \bigcup_{k \in B_i} \overline{N}_{\delta/2}(F_k) \cup \bigcup_{k \in B'_i} \overline{N}_{\delta/2}(F'_k) \right),$$

where $\overline{N}_\epsilon(V)$ denotes the closed ϵ -neighborhood of V . Each \tilde{U}_i is open, and as $\delta \leq \delta_1$, it is clear that $\bigcup_i \tilde{U}_i$ contains $[0, 1]^n$.

We claim that for each i and j , either $\tilde{U}_i \cap \tilde{U}_j \neq \emptyset$ or $\text{dist}(\tilde{U}_i, \tilde{U}_j) \geq \delta/2$. Indeed, if $\text{dist}(\tilde{U}_i, \tilde{U}_j) < \delta/2$, then there are $x \in \tilde{U}_i$ and $y \in \tilde{U}_j$ with $|x - y| < \delta/2$. This implies that $U_i \cap U_j \neq \emptyset$, for if not, then $j \in A_i$ so that x could not be in \tilde{U}_i . Choose $z \in U_i \cap U_j$ with $B(z, \delta) \subset U_i \cap U_j$, which is possible because $\delta \leq \delta_2$. Then it must be that $z \in \tilde{U}_i \cap \tilde{U}_j$. Indeed, suppose that $z \notin \tilde{U}_i$. Then either there is $l \in A_i$ with $z \in \overline{N}_{\delta/2}(U_l)$ or there is $k \in B_i \cup B'_i$ with $z \in \overline{N}_{\delta/2}(F_k) \cup \overline{N}_{\delta/2}(F'_k)$. However, as $B(z, \delta) \subset U_i$, the distance from z to any of these U_l , F_k , or F'_k is strictly larger than $\delta/2$. This immediately rules out $z \in \overline{N}_{\delta/2}(U_l)$ or $z \in \overline{N}_{\delta/2}(F_k) \cup \overline{N}_{\delta/2}(F'_k)$. The argument for $z \in \tilde{U}_j$ is the same.

Similarly, we can also show that for each i and k , either $\tilde{U}_i \cap F_k \neq \emptyset$ or $\text{dist}(\tilde{U}_i, F_k) \geq \delta/2$. Indeed, if $\text{dist}(\tilde{U}_i, F_k) < \delta/2$, then $U_i \cap F_k \neq \emptyset$. This implies that there is $z \in F_k$ for which $B(z, \delta) \subset U_i$. As $z \notin \overline{N}_{\delta/2}(U_j)$ for each $j \in A_i$, we necessarily have $z \in \tilde{U}_i$. Hence, $\tilde{U}_i \cap F_k \neq \emptyset$. The same arguments also show that for each i and k , either $\tilde{U}_i \cap F'_k \neq \emptyset$ or $\text{dist}(\tilde{U}_i, F'_k) \geq \delta/2$. Thus, the collection $\{\tilde{U}_i\}_{i \in I}$ has the convenient property that for every incidence relevant to the calculation of a combinatorial distance, the associated sets either intersect or are of distance $\geq \delta/2$ from each other.

Let us again modify the collection $\{\tilde{U}_i\}_{i \in I}$ slightly, in a way dependent on a parameter $\epsilon > 0$ that we will eventually send to 0. Namely, let $0 < \epsilon < \delta/(8\sqrt{n})$ be very small. For each i , let $V_i = N_{\epsilon/2}(\tilde{U}_i)$, where the neighborhood is now taken in \mathbb{R}^n . Thus, each V_i is open in \mathbb{R}^n , and the union $\bigcup_i V_i$ contains $[0, 1]^n$ but does not intersect any of the half-spaces $\pi_k^{-1}([1 + \epsilon/2, \infty))$.

To the collection $\{V_i\}_{i \in I}$ we add small Euclidean balls to produce a cover of the cube $[0, 1 + \epsilon]^n$. Namely, we can find a collection of points $\{x_j\}_{j \in J}$ with the following properties: $\#J \leq C_n(1/\epsilon)^{n-1}$, where C_n is a dimensional constant; each point x_j lies in one of the codimension-1 spaces $\{x \in \mathbb{R}^n : \pi_k(x) = 1 + \epsilon\}$; and the balls $B_j = B(x_j, \sqrt{n}\epsilon)$ have

$$(1, 1 + \epsilon]^n \subset \bigcup_{j \in J} B_j.$$

Let \mathcal{V} denote the collection of open sets $\{V_i\}_{i \in I} \cup \{B_j\}_{j \in J}$ so that \mathcal{V} is an open cover of the cube $[0, 1 + \epsilon]^n$. Let

$$G_k = \pi_k^{-1}(\{1 + \epsilon\}) \cap [0, 1 + \epsilon]^n$$

be the codimension-1 face of $[0, 1 + \epsilon]^n$ opposite to F_k , and observe that no set in \mathcal{V} intersects both F_k and G_k . In other words, \mathcal{V} is non-spanning (the fact that we are covering a slightly larger cube is not a problem; indeed, Proposition 3.1 applies equally well to topological cubes).

To obtain weight functions for \mathcal{V} , we of course want to use the original weights w_k associated to the cover \mathcal{U} . Namely, let v_k be weight functions for \mathcal{V} defined by

$$v_k(V_i) = w_k(U_i) \quad \text{and} \quad v_k(B_j) = \epsilon,$$

and let $\tilde{d}_k = \text{dist}_{v_k}(F_k, G_k)$ be the associated distance between opposite faces of the cube. We claim that $d_k \leq \tilde{d}_k$ for each k . To see this, let us first establish that any chain in \mathcal{V} of minimal v_k -length that connects F_k to G_k has the form

$$V_{i_1}, \dots, V_{i_m}, B_j$$

for some collection i_1, \dots, i_m and some j . It is clear that the chain must end with some B_j , as none of the V_i intersect G_k . Also, each ball B_j intersects G_k , so the penultimate set in the chain cannot be some other ball B_l . Lastly, note that if $V_i, B_j, V_{i'}$ appears in the chain, then $\text{dist}(V_i, V_{i'}) \leq \text{diam } B_j \leq 2\sqrt{n}\epsilon$, so that

$$\text{dist}(\tilde{U}_i, \tilde{U}_{i'}) \leq 4\sqrt{n}\epsilon < \delta/2.$$

Consequently, $\tilde{U}_i \cap \tilde{U}_{i'} \neq \emptyset$, which also means that $V_i \cap V_{i'} \neq \emptyset$. Thus, in a minimal chain, a ball B_j never appears between two of the V_i 's.

Let $V_{i_1}, \dots, V_{i_m}, B_j$ be a chain of minimal v_k -length from F_k to G_k . As $V_{i_l} \cap V_{i_{l+1}} \neq \emptyset$ for each l , we know that $\text{dist}(\tilde{U}_{i_l}, \tilde{U}_{i_{l+1}}) \leq \epsilon < \delta/2$. Thus, $\tilde{U}_{i_l} \cap \tilde{U}_{i_{l+1}} \neq \emptyset$, so also $U_{i_l} \cap U_{i_{l+1}} \neq \emptyset$. Hence, U_{i_1}, \dots, U_{i_m} is a chain in the collection \mathcal{U} . Moreover, $\text{dist}(\tilde{U}_{i_1}, F_k) \leq \epsilon/2 < \delta/2$, so $\tilde{U}_{i_1} \cap F_k \neq \emptyset$, and also $U_{i_1} \cap F_k \neq \emptyset$. Similarly,

$$\text{dist}(\tilde{U}_{i_m}, F'_k) \leq \epsilon/2 + \text{diam } B_j \leq \epsilon/2 + 2\sqrt{n}\epsilon < \delta/2,$$

so that $\tilde{U}_{i_m} \cap F'_k \neq \emptyset$, and also $U_{i_m} \cap F'_k \neq \emptyset$. Therefore, the chain U_{i_1}, \dots, U_{i_m} connects F_k and F'_k , which implies that

$$d_k \leq \sum_{l=1}^m w_k(U_{i_l}) = \sum_{l=1}^m v_k(V_{i_l}) = \tilde{d}_k - \epsilon \leq \tilde{d}_k.$$

Applying Proposition 3.1 to the collection \mathcal{V} with weight functions v_k gives

$$\prod_{k=1}^n d_k \leq \prod_{k=1}^n \tilde{d}_k \leq \sum_{i \in I} \left(\prod_{k=1}^n v_k(V_i) \right) + \sum_{j \in J} \epsilon^n \leq \sum_{i \in I} \left(\prod_{k=1}^n w_k(U_i) \right) + C_n \epsilon,$$

where the last inequality follows from the bound $\#J \leq C_n(1/\epsilon)^{n-1}$. As this holds for any $0 < \epsilon < \delta/(8\sqrt{n})$, we send ϵ to zero to obtain

$$\prod_{k=1}^n d_k \leq \sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right).$$

□

Remark 3.2. Not surprisingly, the methods we have used to prove Theorem 1.5 can be adapted to prove similar inequalities on other convex polyhedra. In the Riemannian setting, Derrick's methods were extended to a much more general framework by M. Gromov, and this includes a diameter-volume inequality for simplices [7, Section 7]. We will not attempt to build an analogous framework here, but we should remark that modifying the arguments above to establish discrete diameter-volume inequalities for simplices is fairly straightforward. More specifically, let Δ^n denote the n -dimensional simplex, and let T_1, \dots, T_{n+1} denote its codimension-1 faces. If

$\mathcal{U} = \{U_i\}_{i \in I}$ is an open cover of Δ^n and $w: I \rightarrow [0, \infty)$ is a corresponding weight function, then we define the *diameter* of \mathcal{U} , with respect to w , to be

$$d_w(\mathcal{U}) = \inf \left\{ \sum_{j=1}^m w(i_j) : \begin{array}{l} \text{the collection } U_{i_1}, \dots, U_{i_m} \text{ contains chains} \\ \text{that connect } T_k \text{ and } T_l \text{ for each pair } k, l \end{array} \right\}.$$

Adapting the ideas in our proof of Theorem 1.5 to this setting, one can prove that

$$\sum_{i \in I} w(i)^n \geq \frac{d_w(\mathcal{U})^n}{n!}.$$

Moreover if $w \equiv 1$ is constant, then a modification of the arguments used in [13] to prove Theorem 1.3 gives the improved estimate

$$\#\mathcal{U} \geq \binom{n + d_w(\mathcal{U}) - 1}{n}.$$

This is carried out in the author's Ph.D. thesis [12, Theorem 3.17].

4. LOWER VOLUME BOUNDS IN METRIC SPACES

Using Theorem 1.5, we can prove a similar length-volume inequality for *images* of the Euclidean cube in a metric space. To set this up, let (X, d) be a metric space and let $g: [0, 1]^n \rightarrow X$ be a continuous map. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $g([0, 1]^n)$ with corresponding non-negative weight functions w_k for $1 \leq k \leq n$. If $A, B \subset g([0, 1]^n)$ are subsets, then we define

$$\text{dist}_{w_k}(A, B) = \inf \left\{ \sum_{j=1}^m w_k(i_j) : \begin{array}{l} U_{i_1}, \dots, U_{i_m} \text{ is a chain} \\ \text{that connects } A \text{ and } B \end{array} \right\}$$

as before, where *chains* are finite sequences of the $\{U_i\}_{i \in I}$ whose consecutive sets have non-empty intersection. Let

$$d_k(g) = \text{dist}_{w_k}(g(F_k), g(F'_k))$$

be the discrete distance between the images of opposite faces. Of course, this distance depends strongly on \mathcal{U} and w_k as well.

Proposition 4.1. *For such $g: [0, 1]^n \rightarrow X$, we have*

$$\sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right) \geq \prod_{k=1}^n d_k(g).$$

Proof. For each $i \in I$, let $V_i = g^{-1}(U_i)$, so that the collection $\mathcal{V} = \{V_i\}_{i \in I}$ forms an open cover of $[0, 1]^n$. Let $d_k = \text{dist}_{w_k}(F_k, F'_k)$ be the discrete distance associated to the cover \mathcal{V} . Observe that if V_{i_1}, \dots, V_{i_m} is a chain in \mathcal{V} that connects F_k and F'_k , then U_{i_1}, \dots, U_{i_m} is a chain in \mathcal{U} that connects $g(F_k)$ and $g(F'_k)$. Consequently, we have $d_k(g) \leq d_k$ for each k . Applying Theorem 1.5 to the cover \mathcal{V} with weights w_k gives

$$\prod_{k=1}^n d_k(g) \leq \prod_{k=1}^n d_k \leq \sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right),$$

as desired. \square

Using this proposition, we can establish a similar inequality relating more standard metric quantities such as Hausdorff measure and metric distance between sets. Recall from Section 1 that if (X, d) is a metric space and $E \subset X$ is compact, the Q -dimensional Hausdorff content of E is

$$\mathcal{H}_Q^\infty(E) = \inf \left\{ \sum_{i \in I} (\text{diam } U_i)^Q : \{U_i\}_{i \in I} \text{ is an open cover of } E \right\}.$$

The associated Hausdorff Q -dimensional measure is defined to be

$$\mathcal{H}_Q(E) = \liminf_{\epsilon \searrow 0} \left\{ \sum_{i \in I} (\text{diam } U_i)^Q : \{U_i\}_{i \in I} \text{ covers } E \text{ and } \text{diam } U_i < \epsilon \right\},$$

and it is clear that $\mathcal{H}_Q(E) \geq \mathcal{H}_Q^\infty(E)$. Thus, lower bounds on Hausdorff content are also lower bounds on Hausdorff measure.

Corollary 4.2. *If $g: [0, 1]^n \rightarrow X$ is continuous, then*

$$\mathcal{H}_n^\infty(g([0, 1]^n)) \geq \prod_{k=1}^n \text{dist}(g(F_k), g(F'_k)).$$

Proof. Fix an open cover $\{U_i\}_{i \in I}$ of $g([0, 1]^n)$, and let $w_k(i) = \text{diam } U_i$ for each i and each $1 \leq k \leq n$. Observe that if U_{i_1}, \dots, U_{i_m} is a chain connecting $g(F_k)$ and $g(F'_k)$, then

$$\text{dist}(g(F_k), g(F'_k)) \leq \sum_{j=1}^m \text{diam } U_{i_j} = \sum_{j=1}^m w_k(i_j).$$

Thus, $\text{dist}(g(F_k), g(F'_k)) \leq \text{dist}_{w_k}(g(F_k), g(F'_k))$ for each k . By Proposition 4.1, we have

$$\sum_{i \in I} (\text{diam } U_i)^n \geq \prod_{k=1}^n \text{dist}_{w_k}(g(F_k), g(F'_k)) \geq \prod_{k=1}^n \text{dist}(g(F_k), g(F'_k)).$$

As this holds for any open cover $\{U_i\}_{i \in I}$ of $g([0, 1]^n)$, we obtain the desired inequality. \square

Let us note that Corollary 4.2 is sharp in the sense that equality holds when $X = ([0, 1]^n, \ell_\infty)$ and g is the identity map. In this case, $\text{dist}_X(F_k, F'_k) = 1$ for each k , and it is straightforward to show even that $\mathcal{H}_Q(X) = 1$. Indeed, one can cover $[0, 1]^n$ by 2^{jn} cubes of side-length 2^{-j} for each $j \in \mathbb{N}$, and the ℓ_∞ -diameter of such cubes is 2^{-j} . Thus, it is important that we *do not* define Hausdorff content (or Hausdorff measure) with the normalizing factor $\text{Vol}(B^n)2^{-n} < 1$.

We also observe that in both Proposition 4.1 and Corollary 4.2, it is not necessary that X be a metric space. The same arguments hold if X is a *pseudometric space*: the distance between distinct points is allowed to be zero. We illustrate this with the following result, which essentially answers a question of Y. Burago and V. Zalgaller [3, p. 296].

Corollary 4.3. *Let ρ be a pseudometric on $[0, 1]^n$, and assume that every open set in the topology determined by ρ is also open in the Euclidean topology. Then*

$$\mathcal{H}_{n,\rho}^\infty([0, 1]^n) \geq \prod_{k=1}^n \text{dist}_\rho(F_k, F'_k).$$

Here, $\mathcal{H}_{n,\rho}^\infty([0,1]^n)$ and $\text{dist}_\rho(F_k, F'_k)$ are defined in the same manner as the usual Hausdorff content and distance, but using the pseudometric ρ instead of an actual metric. Burago and Zalgaller asked whether the inequality in Corollary 4.3 is true with Hausdorff measure in place of Hausdorff content, though it seems that their definition of Hausdorff measure includes the factor $\text{Vol}(B^n)2^{-n}$. Taken on face-value, this question has a negative answer, as is exhibited by the metric space $([0,1]^n, \ell_\infty)$. Corollary 4.3 appears to be the strongest statement that holds in the general metric setting. It does not, however, immediately recover the inequalities of Derrick and Almgren that are discussed in [3, Section 38.1].

Proof. The pseudometric ρ on $X = [0,1]^n$ canonically induces a metric $\tilde{\rho}$ on the quotient space $\tilde{X} = X/\sim$, where \sim is the equivalence relation $x \sim x'$ if $\rho(x, x') = 0$. Let $\pi: X \rightarrow \tilde{X}$ be the associated projection. It is straightforward to see that

$$(4.1) \quad \text{dist}_\rho(F_k, F'_k) = \text{dist}_{\tilde{\rho}}(\pi(F_k), \pi(F'_k)) \quad \text{and} \quad \mathcal{H}_{n,\rho}^\infty(X) = \mathcal{H}_{n,\tilde{\rho}}^\infty(\tilde{X}).$$

Moreover, $U \subset \tilde{X}$ is open in the metric topology if and only if $\pi^{-1}(U)$ is open in the topology on $[0,1]^n$ determined by ρ . The hypothesis of topological compatibility then ensures that π , when viewed as a map from the Euclidean cube $[0,1]^n$ to the metric space \tilde{X} , is continuous. Corollary 4.2, along with (4.1), gives the desired conclusion. \square

Corollary 4.2 points us in the following direction: in what generality can one obtain Euclidean-type lower volume bounds in metric spaces? More precisely, for which metric spaces (X, d) , does one have

$$(4.2) \quad \mathcal{H}_n^\infty(B(x, r)) \gtrsim r^n$$

for all metric balls $B(x, r)$ with $0 < r \leq \text{diam } X$? An immediate consequence of Corollary 4.2 is that (4.2) holds whenever X satisfies the following property.

Definition 4.4. A metric space (X, d) is said to *admit fat n -cubes* if there is $\lambda \geq 1$ such that, for each $x \in X$ and $0 < r \leq \text{diam } X$, there is a continuous map $g: [0,1]^n \rightarrow B(x, r)$ with $\text{dist}(g(F_k), g(F'_k)) \geq r/\lambda$ for each k .

Euclidean-type lower volume bounds cannot, of course, hold in complete generality: the existence of “thin necks” or “outward cusps” in X would hinder large volume in certain regions. Thus, it makes sense to impose appropriate connectivity conditions when addressing such questions. For $\lambda \geq 1$, we say that X is λ -linearly locally contractible if for each $x \in X$ and $0 < r \leq (\text{diam } X)/\lambda$, the metric ball $B(x, r)$ can be contracted within $B(x, \lambda r)$ to a point. We say that X is linearly locally contractible if it is λ -linearly locally contractible for some value λ . In this context, it turns out that there is a fairly general result on lower volume bounds, which relies on the following deep fact proved by S. Semmes.

Theorem 4.5 (Semmes [20, Theorem 1.29(a)]). *Let (X, d) be a closed manifold of dimension $n \geq 2$ that is N -doubling and λ -linearly locally contractible. Then for each $x \in X$ and each $0 < r \leq \text{diam } X$, there is a surjective map $f: X \rightarrow \mathbb{S}^n$ that is C/r -Lipschitz and is constant outside of $B(x, r/2)$. Here, C depends only on n , N , and λ .*

Here, N -doubling means that every metric ball of radius $r > 0$ can be covered by at most N balls of radius $r/2$. It functions as a finite-dimensionality condition for metric spaces. From Semmes’s theorem, we immediately obtain an estimate of the

form in (4.2). Indeed, as f has non-zero degree, $f(B(x, r)) = \mathbb{S}^n$, so the Lipschitz bounds imply that

$$\mathcal{H}_n^\infty(B(x, r)) \geq \left(\frac{r}{C}\right)^n \mathcal{H}_n^\infty(\mathbb{S}^n) \gtrsim r^n,$$

where the implicit constant depends only on n , N , and λ .

This ‘‘Semmes approach’’ to lower volume bounds is, in a sense, dual to our approach, which seeks to map a nice space *into* X , rather than map X into some other controlled space. The relative ease of building Lipschitz maps from a general metric space into Euclidean spaces (for example, as we did in Section 2) makes the Semmes method viable. Still, it would be desirable to make our dual argument work. In the next section, we carry this out for metric surfaces. The methods we use also allow us to extend [13, Proposition 4.1], which gives conditions under which a metric space is bi-Lipschitz equivalent to a ‘‘snowflake.’’

5. FAT SQUARES IN METRIC SURFACES

Our work in the previous sections was originally motivated by quasiconformal uniformization problems for metric spaces. Here, one attempts to find global parameterizations of metric spaces which enjoy some sort of ‘‘analytic regularity.’’ These problems are especially important to the geometric study of boundaries of hyperbolic groups. In many situations, though, the metric spaces contain no rectifiable curves, so standard notions of length are useless in this analysis.

The theory that has developed around quasiconformal uniformization problems is quite deep and is especially rich for surfaces. The following theorem is characteristic of the subject. It gives conditions under which a topological sphere must actually be a *quasi-sphere*, i.e., must be a quasisymmetric image of the Euclidean sphere.

Theorem 5.1 (Bonk–Kleiner [1]). *Suppose that (X, d) is homeomorphic to \mathbb{S}^2 , is Ahlfors 2-regular, and is linearly locally contractible. Then X is quasisymmetrically equivalent to \mathbb{S}^2 .*

A metric space X is said to be Ahlfors Q -regular if every metric ball $B(x, r)$ with radius $0 < r \leq \text{diam } X$ satisfies the volume estimate $\mathcal{H}_Q(B(x, r)) \approx r^Q$, with a uniform implicit constant. Recall from the previous section that linear local contractibility means that every metric ball can be contracted within a slightly larger ball to a point. As quasisymmetric maps distort relative distances by a controlled amount, they preserve linear local contractibility. Of course, the Euclidean sphere \mathbb{S}^2 is linearly locally contractible, so every quasi-sphere is as well. Thus, the more restrictive hypothesis in the theorem above is that of Ahlfors 2-regularity. Later in this section, we will extend Theorem 5.1 outside of the 2-regular setting.

For now, though, let us discuss connectivity conditions that are related to linear local contractibility. A common theme among them is that they appear in the study of boundaries of hyperbolic groups. We first list them and then discuss how they relate to each other. For a metric space (X, d) and $\lambda \geq 1$, we say that

- (i) X is λ -LLC₁ if, for each $p \in X$ and $0 < r \leq \text{diam } X$, any two points $x, y \in B(p, r)$ can be joined by a continuum in $B(p, \lambda r)$;
- (ii) X is λ -LLC₂ if, for each $p \in X$ and $0 < r \leq \text{diam } X$, any two points $x, y \in X \setminus B(p, r)$ can be joined by a continuum in $X \setminus B(p, r/\lambda)$;
- (iii) X is λ -annularly linearly connected if it is connected and, for each $p \in X$ and $0 < r \leq \text{diam } X$, any two points $x, y \in A(p, r, 2r)$ can be joined by a continuum in $A(p, r/\lambda, 2\lambda r)$.

Here, we use $A(p, r, R) = \overline{B}(p, R) \setminus B(p, r)$ to denote the (closed) metric annulus centered at p with inner radius $r > 0$ and outer radius $R > r$. Recall that a continuum is simply a compact connected set. The “LLC” acronym in LLC_1 and LLC_2 stands for “linearly locally connected” and should not be confused with linear local contractibility. For convenience, we will also use the acronym “ALC” in place of “annularly locally connected.” This third condition is the strongest of the three; it was introduced by J. Mackay in [17].

Lemma 5.2. *If X is λ -ALC, then it is λ' - LLC_1 and λ' - LLC_2 for some λ' , depending only on λ .*

Proof. We first verify the LLC_2 property. Let $p \in X$, let $0 < r \leq \text{diam } X$, and fix $x, y \in X \setminus B(p, r)$. Without loss of generality, suppose that $R = d(p, x) \leq d(p, y)$, and let n be the largest integer for which $2^n R < d(p, y)$. Let $x_0 = x$, and for each $k \leq n$, choose $x_k \in X$ with $d(p, x_k) = 2^k R$. This is possible because X is connected. Finally, let $x_{n+1} = y$. The ALC condition guarantees that for each $0 \leq k \leq n$, there is a continuum

$$E_k \subset A\left(p, \frac{2^k R}{\lambda'}, 2^{k+1} R \lambda'\right)$$

connecting x_k and x_{k+1} , as long as $\lambda' > \lambda$. In particular, the continuum $E = E_0 \cup \dots \cup E_n$ connects x and y and is contained in $X \setminus B(p, r/\lambda')$. Thus, X is λ' - LLC_2 for every $\lambda' > \lambda$.

To verify the LLC_1 condition, fix $x, y \in B(p, r)$. By the connectivity of X , we may choose $q \in X$ for which $d(x, q) = d(x, y)/2$. Then

$$x, y \in A\left(q, \frac{d(x, y)}{4}, \frac{3d(x, y)}{2}\right),$$

and using the same technique as in the previous paragraph, it is not difficult to show that x and y can be connected by a continuum

$$E \subset A\left(q, \frac{d(x, y)}{4\lambda'}, \frac{3\lambda' d(x, y)}{2}\right),$$

where λ' depends only on λ . In particular, $E \subset B(p, 5\lambda' r)$. Thus, we see that X is $5\lambda'$ - LLC_1 . \square

When X has some topological regularity, there are close relationships between the LLC_1 and LLC_2 conditions, linear local contractibility, and the ALC condition.

Lemma 5.3. *For (X, d) a closed, connected manifold of dimension $n \geq 2$, the following are true.*

- (i) *If X is λ -linearly locally contractible, then it is λ' - LLC_1 and λ' - LLC_2 for all $\lambda' > \lambda$.*
- (ii) *If $n = 2$ and X is both LLC_1 and LLC_2 , then it is linearly locally contractible.*
- (iii) *If X is λ -linearly locally contractible, then it is λ' -ALC for all $\lambda' > \lambda$.*

Proof. Parts (i) and (ii) are directly from [1, Lemma 2.5]. Part (iii) is effectively a modification of the proof of (i), along with excision. Fix $p \in X$ and $0 < r \leq \text{diam } X$, and let $\lambda' > \lambda$. As a first case, suppose that $r > \text{diam } X/\lambda$, so $B(p, 2\lambda' r) = X$. Part (i) implies that X is λ' - LLC_2 , so any two points in $A(p, r, 2r)$ can be joined by a continuum in $X \setminus B(p, r/\lambda') = A(p, r/\lambda', 2\lambda' r)$. Thus, we may assume that $r \leq \text{diam } X/\lambda$.

Let us define

$$U = \{x \in X : x \text{ can be joined to } p \text{ by a path in } B(p, 2\lambda'r)\},$$

which is an open and path connected subset of $B(p, 2\lambda'r)$. Part (i) implies that X is λ' -LLC₁ for any $\lambda' > \lambda$, and this ensures that $\overline{B}(p, 2r) \subset U$. It therefore suffices to show that any two points in $U \setminus B(p, r)$ can be joined by a path in $U \setminus B(p, r/\lambda')$. In terms of reduced homology groups, this is equivalent to the homomorphism

$$i_* : \tilde{H}_0(U \setminus B(p, r)) \rightarrow \tilde{H}_0(U \setminus B(p, r/\lambda')),$$

induced by the inclusion map, being identically zero. Here we use singular homology with coefficients in \mathbb{Z}_2 so that there are no issues with orientation.

As $\lambda' > \lambda$, we may choose $0 < r' < r$ such that $\overline{B}(p, r/\lambda') \subset B(p, r'/\lambda)$. Let $K_1 = \overline{B}(p, r/\lambda')$ and $K_2 = \overline{B}(p, r')$ so that both K_1 and K_2 are compact and

$$(5.1) \quad B(p, r/\lambda') \subset K_1 \subset K_2 \subset B(p, r) \subset U.$$

Note, additionally, that $K_1 \subset B(p, r'/\lambda)$, which implies that K_1 can be contracted within K_2 to a point. The containments in (5.1) guarantee that i_* above will be zero if the homomorphism $\tilde{H}_0(U \setminus K_2) \rightarrow \tilde{H}_0(U \setminus K_1)$ induced by inclusion is zero.

The long exact sequence for relative homology [9, p. 117], together with path connectedness of U , gives that

$$\dots \rightarrow H_1(U, U \setminus K_i) \xrightarrow{\partial_i} \tilde{H}_0(U \setminus K_i) \rightarrow \tilde{H}_0(U) = 0$$

is exact for $i = 1, 2$. Thus, ∂_i is surjective. These sequences are natural with respect to inclusions, so triviality of $\tilde{H}_0(U \setminus K_2) \rightarrow \tilde{H}_0(U \setminus K_1)$ would be guaranteed by triviality of the homomorphism $H_1(U, U \setminus K_2) \rightarrow H_1(U, U \setminus K_1)$, also induced by inclusion.

Applying excision [9, p. 119] to the set $X \setminus U$, which is compact and contained in $X \setminus K_i$, we find that the inclusion $(U, U \setminus K_i) \hookrightarrow (X, X \setminus K_i)$ induces an isomorphism $H_1(U, U \setminus K_i) \rightarrow H_1(X, X \setminus K_i)$. It is not difficult to see that this isomorphism is natural with respect to inclusions (indeed, it is induced by an inclusion!), so the following diagram commutes:

$$\begin{array}{ccc} H_1(U, U \setminus K_2) & \rightarrow & H_1(X, X \setminus K_2) \\ \downarrow & & \downarrow \\ H_1(U, U \setminus K_1) & \rightarrow & H_1(X, X \setminus K_1) \end{array}$$

where the vertical maps come from the inclusions $(U, U \setminus K_2) \hookrightarrow (U, U \setminus K_1)$ and $(X, X \setminus K_2) \hookrightarrow (X, X \setminus K_1)$. Thus, in order to show that the vertical map on the left is the zero homomorphism, it suffices to prove this for the vertical map on the right. Our goal, then, is to show that $H_1(X, X \setminus K_2) \rightarrow H_1(X, X \setminus K_1)$ is trivial.

For this, we use Poincaré duality [22, p. 296]. As K_i is compact, there is an isomorphism $H_1(X, X \setminus K_i) \simeq \check{H}^{n-1}(K_i)$, where \check{H}^k denotes the k -th Čech cohomology group with \mathbb{Z}_2 coefficients. Once again, this isomorphism is natural with respect to inclusions, so triviality of $H_1(X, X \setminus K_2) \rightarrow H_1(X, X \setminus K_1)$ is equivalent to triviality of the homomorphism $\check{H}^{n-1}(K_2) \rightarrow \check{H}^{n-1}(K_1)$ induced by the inclusion $K_1 \hookrightarrow K_2$. Recall, though, that K_1 is contractible inside K_2 , so this latter homomorphism is indeed trivial.

We conclude, then, that the original homomorphism i_* is identically zero, which means that any two points in $U \setminus B(p, r)$ can be joined by a path in $U \setminus B(p, r/\lambda')$. As $\overline{B}(p, 2r) \subset U \subset B(p, 2\lambda'r)$, this implies that X is λ' -ALC. \square

A major motivation for introducing these types of connectivity conditions is that they appear in the analysis of hyperbolic groups. Namely, if G is a hyperbolic group and $\partial_\infty G$ denotes its boundary at infinity equipped with a visual metric, then under suitable topological hypotheses, $\partial_\infty G$ will satisfy all of these conditions. For example, if $\partial_\infty G$ is non-empty, connected, and has no local cut points (equivalently, G does not split over a finite group or over a virtually cyclic group), then it is ALC [17, Proof of Corollary 1.2]. Similarly, if $\partial_\infty G$ is a connected manifold, then it is linearly locally contractible ([11, Theorem 4.4], along with [14, Theorem 3.3]).

Recall from the previous section that (X, d) is said to admit fat n -cubes if there is $\lambda \geq 1$ such that, for each $x \in X$ and $0 < r \leq \text{diam } X$, there is a continuous map $g: [0, 1]^n \rightarrow B(x, r)$ with $\text{dist}(g(F_k), g(F'_k)) \geq r/\lambda$ for each k . The following condition is slightly stronger and is related to the connectivity conditions above.

Definition 5.4. We say that (X, d) *admits fat connecting n -cubes* if there is $\lambda \geq 1$ such that for any two distinct points $x, y \in X$ with $d(x, y) \leq (\text{diam } X)/\lambda$, there is a continuous map $g: [0, 1]^n \rightarrow B(x, \lambda d(x, y))$ with $g(F_1) \subset B(x, d(x, y)/4)$ and $g(F'_1) \subset B(y, d(x, y)/4)$, and which has $\text{dist}(g(F_k), g(F'_k)) \geq d(x, y)/\lambda$ for each $1 \leq k \leq n$.

When $n = 2$, we use the terminology “fat connecting squares” in place of “fat connecting 2-cubes”. The following theorem of J. Mackay will be helpful in comparing the “fat squares” condition to the other connectivity conditions. Recall that X is said to be N -doubling if every ball of radius $r > 0$ can be covered by at most N balls of radius $r/2$. An arc γ in X is called an α -quasiarc if for each pair $x, y \in \gamma$, the sub-arc between them has diameter at most $\alpha d(x, y)$. Similarly, a topological circle in X is called an α -quasicircle if for each pair of points x, y on it, there is a sub-arc between them with diameter at most $\alpha d(x, y)$.

Theorem 5.5 (Mackay [16, Theorem 1.4]). *Suppose that (X, d) is a complete, N -doubling, λ -ALC metric space. For any $m \in \mathbb{N}$, there is $\alpha = \alpha(N, \lambda, m) \geq 1$ such that any two distinct points $x, y \in X$ can be joined by m different α -quasiarcs, such that the concatenation of any two of them is an α -quasicircle.*

The following lemma tells us that linear local contractibility and annular linear connectivity give fat connecting squares.

Lemma 5.6. *Let (X, d) be a complete metric space that is N -doubling, λ -linearly locally contractible, and λ' -ALC. Then X admits fat connecting squares with constant depending only on N , λ , and λ' .*

Proof. Fix $x, y \in X$ distinct, and for notational ease let $r = d(x, y) > 0$. By Theorem 5.5, there is $\alpha \geq 1$, depending only on N and λ' , such that x and y can be joined by two α -quasiarcs in X for which the concatenation of the two is an α -quasicircle. Let \mathcal{C} denote this quasicircle, and observe that $\mathcal{C} \subset B(x, \alpha r)$.

Let γ_1 be a sub-arc of \mathcal{C} with $\gamma_1 \subset B(x, r/4)$ and $\text{diam } \gamma_1 \geq r/8$. Similarly, let γ'_1 be a sub-arc of \mathcal{C} with $\gamma'_1 \subset B(y, r/4)$ and $\text{diam } \gamma'_1 \geq r/8$. Then, let γ_2 and γ'_2 be the sub-arcs of \mathcal{C} that connect γ_1 to γ'_1 . Observe that by construction, $\text{dist}(\gamma_1, \gamma'_1) \geq r/2$, and by the quasicircle condition,

$$\text{dist}(\gamma_2, \gamma'_2) \geq \frac{1}{\alpha} \min\{\text{diam } \gamma_1, \text{diam } \gamma'_1\} \geq \frac{r}{8\alpha}.$$

Suppose that $r = d(x, y) \leq (\text{diam } X)/(\alpha\lambda)$. As (X, d) is λ -linearly locally contractible, the ball $B(x, \alpha r)$ is contractible inside $B(x, \alpha\lambda r)$. Thus, there is a continuous map

$$H: [0, 1] \times B(x, \alpha r) \rightarrow B(x, \alpha\lambda r)$$

for which $H_0 = \text{id}$ and $H_1 \equiv \text{const}$; here, we use the standard notation $H_t(z) = H(t, z)$ for $0 \leq t \leq 1$ and $z \in B(x, \alpha r)$. We may restrict this homotopy to the quasicircle, and pre-composing with a parameterization of \mathcal{C} , we obtain

$$\tilde{H}: [0, 1] \times \mathbb{S}^1 \rightarrow B(x, \alpha\lambda r),$$

where $\tilde{H}_0(\mathbb{S}^1) = \mathcal{C}$ and $\tilde{H}_1 \equiv \text{const}$. This defines a continuous map $\tilde{g}: \overline{\mathbb{D}} \rightarrow B(x, \alpha\lambda r)$, where \mathbb{D} denotes the unit disk in \mathbb{R}^2 , and $\tilde{g}|_{\mathbb{S}^1} = \tilde{H}_0|_{\mathbb{S}^1}$ gives a parameterization of \mathcal{C} . More precisely,

$$\tilde{g}(se^{i\theta}) = \tilde{H}(1 - s, e^{i\theta})$$

for each $0 \leq s \leq 1$ and $0 \leq \theta \leq 2\pi$. Notice that γ_1 , γ'_1 , γ_2 , and γ'_2 correspond to sub-arcs of \mathbb{S}^1 under this parameterization. It is straightforward to see that we may pre-compose \tilde{g} with an appropriate homeomorphism from $[0, 1]^2$ to $\overline{\mathbb{D}}$ to obtain a continuous map

$$g: [0, 1]^2 \rightarrow B(x, \alpha\lambda r),$$

where $g(F_1) = \gamma_1$, $g(F'_1) = \gamma'_1$, $g(F_2) = \gamma_2$, and $g(F'_2) = \gamma'_2$. Thus, the map g satisfies $g(F_1) \subset B(x, r/4)$ and $g(F'_1) \subset B(y, r/4)$, along with the estimate $\text{dist}(g(F_2), g(F'_2)) \geq r/(8\alpha)$. This shows that X admits fat connecting squares with constant $\max\{8\alpha, \lambda\alpha\}$. \square

Using Lemma 5.3(iii) and Corollary 4.2, we obtain the following result.

Corollary 5.7. *Let (X, d) be a closed, connected manifold of dimension ≥ 2 that is N -doubling and λ -linearly locally contractible. Then X admits fat connecting squares with constant depending only on N and λ . In particular, $\mathcal{H}_2^\infty(B(x, r)) \gtrsim r^2$ for all balls with $0 < r \leq \text{diam } X$, where the implicit constant again depends only on N and λ .*

The conclusion in this corollary recovers, for surfaces, the statement about lower volume bounds that we made in the previous section, based off of Semmes's theorem. A natural question is whether our method can be extended to higher dimensions. We leave this as an open problem.

5.1. An application to quasi-spheres. Let us return to the quasimetric uniformization result of Bonk and Kleiner in Theorem 5.1. Earlier, those same authors had proved a uniformization theorem that holds in all dimensions, albeit with a strong dynamical hypothesis on the metric space.

Theorem 5.8 (Bonk–Kleiner [2]). *Let (X, d) be a compact, Ahlfors n -regular metric space with topological dimension $n \geq 1$, and suppose that X admits a uniformly quasi-Möbius group action that is cocompact on triples. Then X is quasimetrically equivalent to \mathbb{S}^n .*

Let us not discuss quasi-Möbius group actions here but instead mention that the hypothesis imposes strong self-similarity properties on X . Comparing Theorem 5.1 to the $n = 2$ case of Theorem 5.8, we see that for surfaces, one can essentially drop the assumption of self-similarity, replacing it by the much weaker linear local

contractibility condition, and still obtain the same conclusion regarding quasimetric uniformization.

One might be tempted to ask whether this phenomenon occurs in all dimensions. Namely, if (X, d) is homeomorphic to \mathbb{S}^n , is linearly locally contractible, and is Ahlfors n -regular, is it necessarily quasimetric to the Euclidean sphere? The answer is “no” for $n \geq 3$, by counterexamples due to S. Semmes [21], even though such spaces have good analytic properties [20]. See the Introduction of [1] for further discussion.

In [13], the present author built off of the methods in [2] to obtain a statement similar to that in Theorem 5.8 for metric spaces X whose topological dimension and Ahlfors-regular dimension do not coincide. More specifically, this result gives conditions under which a metric space is bi-Lipschitz equivalent to the “snowflake” of a metric space with equivalent topological and Ahlfors-regular dimensions.

Theorem 5.9 ([13, Proposition 4.1]). *Let $n \geq 2$ and $0 < \epsilon < 1$, and suppose that the metric space (X, d) has the following properties:*

- (i) *X is homeomorphic to \mathbb{S}^n ;*
- (ii) *X admits a conformal elevator;*
- (iii) *every δ -separated set in X has size at most $C\delta^{-n/\epsilon}$;*
- (iv) *every discrete δ -path from x to y has length at least $C^{-1}(d(x, y)/\delta)^{1/\epsilon}$.*

Then there is a metric ρ on X satisfying $\rho \approx d^{1/\epsilon}$, where the implicit constant depends only on the constants from the hypotheses.

Here, a “ δ -separated set” is simply a set of points for which pairwise distances are at least δ . Also, by a “discrete δ -path from x to y ” we mean a chain of points $x = z_0, z_1, \dots, z_m = y$ in X with $d(z_i, z_{i-1}) \leq \delta$ for each i . The length of this chain is m , which is one less than the number of points in the chain. The hypothesis of “admitting a conformal elevator” is a dynamical condition, akin to the assumption that X admits a uniformly quasi-Möbius group action that is cocompact on triples. For general metric spaces, this is a very strong hypothesis. Thus, in the same spirit that motivated Theorem 5.1, it makes sense to ask whether this assumption is necessary.

A close inspection of the proof of Theorem 5.9 in [13] reveals that the first two hypotheses are used *precisely to verify* that X admits fat connecting n -cubes. Thus, we may replace (i) and (ii) by this connectivity assumption. Once this is done, it is also not difficult to modify that proof to deal with metric spaces that are possibly unbounded. Putting these observations together gives the following result.

Proposition 5.10. *Let $n \geq 2$ and $0 < \epsilon < 1$, and suppose that (X, d) is a metric space for which*

- (i') *X admits fat connecting n -cubes with constant λ ,*
- (ii') *every ball of radius $0 < R \leq \text{diam } X$ can be covered by at most $C(R/r)^{n/\epsilon}$ balls of radius $0 < r < R$,*
- (iii') *every discrete δ -path from x to y has length at least $C^{-1}(d(x, y)/\delta)^{1/\epsilon}$.*

Then there is a metric ρ on X satisfying $\rho \approx d^{1/\epsilon}$, where the implicit constant depends only on ϵ , λ , and C .

The condition in (ii') implies that X has Assouad dimension at most $2/\epsilon$. It is not difficult to show that if X is compact and is Ahlfors $2/\epsilon$ -regular, then it automatically satisfies this hypothesis. Specializing the previous proposition to

metric surfaces and using Corollary 5.7, we obtain the following corollary, which builds off of Theorem 5.1.

Corollary 5.11. *Let (X, d) be a closed metric surface that is linearly locally contractible. Suppose that X is Ahlfors $2/\epsilon$ -regular and that every discrete δ -path from x to y in X has length at least $C^{-1}(d(x, y)/\delta)^{1/\epsilon}$. Then there is a metric ρ on X for which $\rho \approx d^{1/\epsilon}$. If, in addition, X is homeomorphic to \mathbb{S}^2 , then X is a quasi-sphere.*

The final statement in this corollary uses the facts that (X, ρ) is Ahlfors 2-regular and that the identity map between (X, d) and (X, ρ) is quasimetric.

REFERENCES

- [1] M. Bonk and B. Kleiner, *Quasisymmetric parametrizations of two-dimensional metric spheres*, Invent. Math. **150** (2002), no. 1, 127–183.
- [2] ———, *Rigidity for quasi-Möbius group actions*, J. Differential Geom. **61** (2002), no. 1, 81–106.
- [3] Y. D. Burago and V. A. Zalgaller, *Geometric inequalities*, Springer-Verlag, Berlin, 1988.
- [4] W. R. Derrick, *Inequalities for p -modules of curve families on Lipschitz surfaces*, Math. Z. **119** (1971), 1–10.
- [5] ———, *A volume-diameter inequality for n -cubes*, J. Analyse Math. **22** (1969), 1–36.
- [6] ———, *A weighted volume-diameter inequality for N -cubes*, J. Math. Mech. **18** (1968/1969), 453–472.
- [7] M. Gromov, *Filling Riemannian manifolds*, J. Differential Geom. **18** (1983), no. 1, 1–147.
- [8] ———, *Metric structures for Riemannian and non-Riemannian spaces*, Progress in Mathematics, vol. 152, Birkhäuser Boston, Boston, MA, 1999.
- [9] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002.
- [10] J. Heinonen, *Lectures on Lipschitz analysis*, Report. University of Jyväskylä Department of Mathematics and Statistics, vol. 100, University of Jyväskylä, Jyväskylä, 2005.
- [11] I. Kapovich and N. Benakli, *Boundaries of hyperbolic groups*, Combinatorial and geometric group theory, Contemp. Math., vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 39–93.
- [12] K. Kinneberg, *A coarse entropy-rigidity theorem and discrete length-volume inequalities*, ProQuest LLC, Ann Arbor, MI, 2014. Thesis (Ph.D.)–UCLA.
- [13] ———, *Rigidity for quasi-Möbius actions on fractal metric spaces*, J. Differential Geom. **100** (2015), no. 2, 349–388.
- [14] B. Kleiner, *The asymptotic geometry of negatively curved spaces: uniformization, geometrization and rigidity*, Proceedings Internat. Congress Math. (Madrid, 2006), Eur. Math. Soc., Zürich, 2006, pp. 743–768.
- [15] J. M. Mackay, *Existence of quasi-arcs*, Proc. Amer. Math. Soc. **136** (2008), no. 11, 3975–3981.
- [16] ———, *Quasi-circles through prescribed points*, Indiana Univ. Math. J. **63** (2014), no. 2, 403–417.
- [17] ———, *Spaces and groups with conformal dimension greater than one*, Duke Math. J. **153** (2010), no. 2, 211–227.
- [18] P. M. Pu, *Some inequalities in certain nonorientable Riemannian manifolds*, Pacific J. Math. **2** (1952), 55–71.
- [19] O. Schramm, *Square tilings with prescribed combinatorics*, Israel J. Math. **84** (1993), no. 1–2, 97–118.
- [20] S. Semmes, *Finding curves on general spaces through quantitative topology, with applications to Sobolev and Poincaré inequalities*, Selecta Math. (N.S.) **2** (1996), no. 2, 155–295.
- [21] ———, *Good metric spaces without good parameterizations*, Rev. Mat. Iberoamericana **12** (1996), no. 1, 187–275.
- [22] E. Spanier, *Algebraic topology*, Springer-Verlag, New York-Berlin, 1981. Corrected reprint.
- [23] K. Wildrick, *Quasisymmetric structures on surfaces*, Trans. Amer. Math. Soc. **362** (2010), no. 2, 623–659.